Hodge decomposition

Let P be an $\mathrm{SO}(n)$ principal bundle. Let $V := P \times_{\rho_{\mathrm{st}}} \mathbb{R}^n$, so that $P = \mathrm{Fr}^{\mathrm{SO}}(V)$. Then we have a fiberwise map $\star : \Lambda^p V \to \Lambda^{n-p} V$ which satisfies

$$\begin{aligned} \alpha \wedge \star\beta &= \langle \alpha, \beta \rangle \; e^1 \wedge \cdots \wedge e^n, \\ \star^2 &= (-1)^{p(n-p)} : \Lambda^p V \to \Lambda^p V. \end{aligned}$$

Suppose X is a smooth n-manifold equipped with a reduction of the cotangent bundle T^*X to a $\mathrm{SO}(n)$ structure, i.e. X is oriented Riemannian. (A Riemannian metric determines an isomorphism $TX \to T^*X$, so reductions of T^*X or TX are equivalent.) In particular, $\star$ induces a bundle map $\Lambda^p T^*X \to \Lambda^{n-p} T^*X$. Differential forms are sections $\Omega^p(X) = \Gamma(\Lambda^p T^*X)$, so we get a map $\star : \Omega^p(X) \to \Omega^{n-p}(X)$ which acts fiberwise.

We define a Euclidean inner product on $\Omega^p_c(X)$ (p-forms with compact supports) by

$$\langle \alpha, \beta \rangle := \int_X \alpha \wedge \star\beta.$$

Define

$$\begin{aligned} d^* &: \Omega^p(X) \to \Omega^{p-1}(X), \\ d^*\alpha &:= (-1)^{n(p+1)+1} \star d \star \alpha. \end{aligned}$$

This satisfies $(d^*)^2 = 0$ as a consequence of $d^2 = 0$ and $\star^2 = \pm 1$.

Theorem. *The operator d^* is the formal metric adjoint of d, i.e. up to a boundary term,*

$$\langle d\alpha, \beta \rangle = \langle \alpha, d^*\beta \rangle .$$

Proof. Suppose $\alpha \in \Omega^{p-1}(X)$ and $\beta \in \Omega^p(X)$. Then

$$d(\alpha \wedge \star\beta) = d\alpha \wedge \star\beta + (-1)^{p-1} \alpha \wedge d \star \beta = d\alpha \wedge \star\beta - \alpha \wedge \star d^*\beta.$$

Integrating, we obtain

$$\int_{\partial X} \alpha \wedge \star\beta = \langle d\alpha, \beta \rangle - \langle \alpha, d^*\beta \rangle .$$

Whenever the boundary term vanishes (e.g. if X is closed), we have $\langle d\alpha, \beta \rangle = \langle \alpha, d^*\beta \rangle$. □

Now let's return to representing cohomology classes on a closed manifold. To find a slice for $\ker d / \mathrm{image}\, d$, we want to consider

$$(\mathrm{image}\; d)^\perp \subset \ker d \subset \Omega^p(X).$$

We have that

$$\alpha \in (\mathrm{image}\; d)^\perp \iff \forall \beta,\; 0 = \langle \alpha, d\beta \rangle = \langle d^*\alpha, \beta \rangle \iff \alpha \in \ker d^*.$$

Since we want to look at the kernel of d^* inside the kernel of d, we are led to study

$$\mathcal{H}^p(X) := \ker d \cap \ker d^* \subset \Omega^p(X).$$

There are a few alternative characterizations. Note that $\mathcal{H}^p(X) = \ker(d \oplus d^*) = \ker(d + d^*)$, since $d + d^* : \Omega^p \to \Omega^{p+1} \oplus \Omega^{p-1}$. For the other characterization, define the Hodge Laplacian

$$\Delta := (d + d^*)^2 = \cancel{d^2} + dd^* + d^*d + \cancel{(d^*)^2}.$$

Over $\mathbb{R}^n$ with the standard metric, Δ on $\Omega^0(X)$ is given by

$$\Delta = -\sum_{i=1}^{n} \left(\frac{\partial}{\partial x^i} \right)^2 .$$

(The minus sign is the geometer's convention, which makes Δ act positively on $e^{i\xi \cdot x} \mapsto \xi^2 e^{i\xi \cdot x}$.)

Note that Δ is formally self-adjoint, since up to boundary terms,

$$\langle \alpha, \Delta\beta \rangle = \langle \alpha, (d + d^*)(d + d^*)\beta \rangle = \langle (d^* + d)\alpha, (d + d^*)\beta \rangle = \langle \Delta\alpha, \beta \rangle .$$

Clearly

$$\alpha \in \mathcal{H}^p(X) \implies (d + d^*)\alpha = 0 \implies \Delta\alpha = 0.$$

But conversely,

$$\Delta\alpha = 0 \implies \langle \alpha, \Delta\alpha \rangle = 0 \implies \langle (d + d^*)\alpha, (d + d^*)\alpha \rangle = 0 \implies (d + d^*)\alpha = 0 \implies \alpha \in \mathcal{H}^p(X).$$

For this reason, $\mathcal{H}^p(X)$ is called the *space of harmonic p-forms*.

We wish to form the decomposition

$$\Omega^p(X) = (\text{image } \Delta) \oplus (\text{image } \Delta)^\perp = (\text{image } \Delta) \oplus (\ker \Delta^*) = (\text{image } \Delta) \oplus \mathcal{H}^p(X).$$

We can further decompose image $\Delta \subset$ (image d) + (image d^*). It is a simple exercise to verify that the images of d and d^* are orthogonal, so we obtain the orthogonal decomposition

$$\Omega^p(X) = d\Omega^{p-1}(X) \oplus d^*\Omega^{p+1}(X) \oplus \mathcal{H}^p(X)$$

known as the Hodge decomposition.

Unfortunately the previous argument was not rigorous, and requires substantial effort. What's wrong with it?

To understand the flaw, consider $V = C^0([0,1])$ the space of continuous functions on the unit interval. Consider the proper subspace $W = C^1([0,1]) \subsetneq V$ of continuously differentiable functions. Unfortunately,

$$V \neq W \oplus W^\perp.$$

In fact, $W^\perp = V^\perp = 0$, since W is L^2-dense in V. Specifically, assume $f \in W^\perp$. We can compute $\langle v, f \rangle = 0$ for any $v \in V$ by taking a sequence of differentiable approximations $w_i \xrightarrow{L^2} v$ which converge to v in the L^2 topology. Then

$$\langle v, f \rangle = \left\langle \lim_{i\to\infty} w_i, f \right\rangle = \lim_{i\to\infty} \langle w_i, f \rangle = \lim_{i\to\infty} 0 = 0.$$

The problem is that W is not closed inside of V in the relative L^2 topology. In order to establish

$$\Omega^p(X) = (\text{image } \Delta) \oplus (\text{image } \Delta)^\perp,$$

we must rule out this sort of possibility.

Let's examine the case $W = C^1([0,1]) \subsetneq C^0([0,1]) = V$ more carefully. Note that V and W are Banach spaces (vector space with a norm, complete with respect to norm-convergence). The C^0 norm on V is given by

$$\|f\|_{C^0} := \sup_x |f(x)|,$$

and the C^1 norm on W is given by

$$\|f\|_{C^1} := \|f\|_{C^0} + \|f'\|_{C^0}.$$

These norms are completely irrelevant (at the moment)! The only relevant norm here is L^2. In particular,

$$W \subset V \subset L^2([0,1]),$$

and the topology of L^2 convergence induces the corresponding subspace topologies on V and W.

More abstractly, consider subspaces $W \subset V \subset H$ of some Hilbert space H. We want to know when $V = W \oplus W^\perp$, where $W^\perp$ denotes the subspace of V orthogonal to W.

Theorem. *$V = W \oplus W^\perp$ iff W is relatively L^2-closed in V.*

Proof. To see necessity of W being relatively L^2-closed, let $\overline{W}_H \subset H$ be the Hilbert space completion of W, so that $V \cap \overline{W}_H =: \overline{W}$ is the relative closure. Then $W \oplus W^\perp \subset \overline{W} \oplus W^\perp \subset V$, and if W is not closed, then the first inclusion would be proper.

To understand why this is sufficient, we can take an orthonormal basis $\{e_i\}$ of $\overline{W}_H$, consisting of vectors in W. The projection to $\overline{W}_H$ is thus $\pi_{\overline{W}_H}(x) = \sum \langle x, e_i \rangle e_i$. If $v \in V$, then $v = \underbrace{\pi_{\overline{W}_H}(v)}_{\overline{W}} + \underbrace{(v - \pi_{\overline{W}_H}(v))}_{W^\perp}$ yields the desired orthonormal decomposition. □

Thus for

$$\Omega^p(X) = (\text{image } \Delta) \oplus (\text{image } \Delta)^\perp,$$

it suffices to show that the image of Δ is relatively L^2-closed.

By the same argument as before with $L = d$, a linear operator L satisfies

$$(\text{image } L)^\perp = \ker L^*.$$

Taking $\perp$ again, we find that

$$\overline{\text{image } L} = (\ker L^*)^\perp.$$

Thus to show that the image of Δ is relatively L^2-closed, it suffices to show that $\mathcal{H}^p(X)^\perp$ is the image of Δ. In other words, if $\alpha \in \Omega^p(X)$ with α orthogonal to $\mathcal{H}^p(X)$, then we need to find $\omega \in \Omega^p(X)$ which solves the PDE $\Delta\omega = \alpha$.

Towards this goal, let X be a closed oriented manifold. Define the L^2 norm of ω by

$$\|\omega\|_{L^2}^2 := \langle \omega, \omega \rangle = \int_X \omega \wedge \star\omega.$$

Lemma 1 (Rellich's Theorem). *Let X be a closed oriented manifold. Suppose $\{\omega_i \in \Omega^p(X)\}_{i=1}^\infty$ is a sequence such that $\|\omega_i\|_{L^2}$ and $\|\Delta\omega_i\|_{L^2}$ are both bounded. Then ω_i contains a Cauchy subsequence, which converges in the L^2-completion of $\Omega^p(X)$, i.e.*

$$\omega_i \xrightarrow{L^2} \omega \in L^2(X; \Lambda^p(T^*X)).$$

This lemma is a fairly deep analytic fact, which we will postpone for later.

Corollary. *If X is a closed oriented manifold, then $\mathcal{H}^p(X)$ is finite-dimensional.*

Proof. Suppose by contradiction that $\mathcal{H}^p(X)$ were infinite-dimensional. Then we could find an infinite orthonormal basis $\{e_i\}$ of $\mathcal{H}^p(X)$. Then $\|e_i\|_{L^2} = 1$ and $\|\Delta e_i\|_{L^2} = 0$ are both bounded, so $\{e_i\}$ defines a Cauchy sequence, which is absurd. $\square$